

Simultaneous Visibility Representations of Plane *st*-graphs Using L-shapes ^{*}

William S. Evans¹, Giuseppe Liotta², and Fabrizio Montecchiani²

¹ University of British Columbia, Canada
will@cs.ubc.ca

² Università degli Studi di Perugia, Italy
{giuseppe.liotta, fabrizio.montecchiani}@unipg.it

Abstract. Let $\langle G_r, G_b \rangle$ be a pair of plane *st*-graphs with the same vertex set V . A simultaneous visibility representation with L-shapes of $\langle G_r, G_b \rangle$ is a pair of bar visibility representations $\langle \Gamma_r, \Gamma_b \rangle$ such that, for every vertex $v \in V$, $\Gamma_r(v)$ and $\Gamma_b(v)$ are a horizontal and a vertical segment, which share an end-point. In other words, every vertex is drawn as an L-shape, every edge of G_r is a vertical visibility segment, and every edge of G_b is a horizontal visibility segment. Also, no two L-shapes intersect each other. An L-shape has four possible rotations, and we assume that each vertex is given a rotation for its L-shape as part of the input. Our main results are: (i) a characterization of those pairs of plane *st*-graphs admitting such a representation, (ii) a cubic time algorithm to recognize them, and (iii) a linear time drawing algorithm if the test is positive.

1 Introduction

Let G_r and G_b be two plane graphs with the same vertex set. A *simultaneous embedding* (SE) of $\langle G_r, G_b \rangle$ consists of two planar drawings, Γ_r of G_r and Γ_b of G_b , such that every edge is a simple Jordan arc, and every vertex is the same point both in Γ_r and in Γ_b . The problem of computing SEs has received a lot of attention in the Graph Drawing literature, partly for its theoretical interest and partly for its application to the visual analysis of dynamically changing networks on a common (sub)set of vertices. For example, it is known that any two plane graphs with the same vertex set admit a SE where the edges are polylines with at most two bends, which are sometimes necessary [9]. If the edges are straight-line segments, the representation is called a *simultaneous geometric embedding* (SGE), and many graph pairs do not have an SGE: a tree and a path [1], a planar graph and a matching [6], and three paths [5]. On the positive side, the discovery of graph pairs that have an SGE is still a fertile research topic. The reader can refer to the survey by Bläsius, Kobourov and Rutter [21] for references and open problems.

Only a few papers study simultaneous representations that adopt a drawing paradigm different from SE and SGE. A seminal paper by Jampani and Lubiw initiates the study of *simultaneous intersection representations* (SIR) [16]. In an intersection representation of a graph, each vertex is a geometric object and there is an edge between two

^{*} Research supported in part by the MIUR project AMANDA “Algorithmics for MASSive and Networked DATA”, prot. 2012C4E3KT_001 and NSERC Canada.

vertices if and only if the corresponding objects intersect. Let $\langle G_r, G_b \rangle$ be two graphs that have a subgraph in common. A SIR of $\langle G_r, G_b \rangle$ is a pair of intersection representations where each vertex in $G_r \cap G_b$ is mapped to the same object in both realizations. Polynomial-time algorithms for testing the existence of SIRs for chordal, comparability, interval, and permutation graphs have been presented [4,15,16].

We introduce and study a different type of simultaneous representation, where each graph is realized as a *bar visibility representation* and two segments representing the same vertex share an end-point. A bar visibility representation of a plane graph G is an embedding preserving drawing Γ where the vertices of G are non-overlapping horizontal segments, and two segments are joined by a vertical visibility segment if and only if there exists an edge in G between the two corresponding vertices (see, e.g., [18,22]). A visibility segment has thickness $\epsilon > 0$ and does not intersect any other segment.

A *simultaneous visibility representation with L-shapes* of $\langle G_r, G_b \rangle$ is a pair of bar visibility representations $\langle \Gamma_r, \Gamma_b \rangle$ such that for every vertex $v \in V$, $\Gamma_r(v)$ and $\Gamma_b(v)$ are a horizontal and a vertical segment that share an end-point. In other words, every vertex is an L-shape, and every edge of G_r (resp., G_b) is a vertical (resp., horizontal) visibility segment. Also, no two L-shapes intersect. A simultaneous visibility representation with L-shapes of $\langle G_r, G_b \rangle$ where the rotation of the L-shape of each vertex in V is defined by a function $\Phi : V \rightarrow \mathcal{H} = \{\lfloor, \lrcorner, \top, \Gamma\}$, is called a Φ -LSVR in the following. While this definition does not assume any particular direction on the edges of G_r (resp., G_b), the resulting representation does induce a bottom-to-top (resp., left-to-right) *st*-orientation. In this paper, we assume that G_r and G_b are directed and this direction must be preserved in the visibility representation. Also, the two graphs have been augmented with distinct (dummy) sources and sinks. More formally, $G_r = (V \cup \{s_r, t_r\}, E_r)$ and $G_b = (V \cup \{s_b, t_b\}, E_b)$ are two plane *st*-graphs with sources s_r, s_b , and sinks t_r, t_b .

In terms of readability, this kind of simultaneous representation has the following advantages: (i) The edges are depicted as straight-line segments (as in SGE) and the edge-crossings are rectilinear; (ii) The edges of the two graphs are easy to distinguish, since they consistently flow from bottom to top for one graph and from left to right for the other graph. Having rectilinear crossing edges is an important benefit in terms of readability, as shown in [14], which motivated a relevant amount of research on right-angle crossing (RAC) drawings, see [10] for a survey.

Our main contribution is summarized by the following theorem.

Theorem 1. *Let G_r and G_b be two plane *st*-graphs defined on the same set of n vertices V and with distinct sources and sinks. Let $\Phi : V \rightarrow \mathcal{H} = \{\lfloor, \lrcorner, \top, \Gamma\}$. There exists an $O(n^3)$ -time algorithm to test whether $\langle G_r, G_b \rangle$ admits a Φ -LSVR. Also, in the positive case, a Φ -LSVR can be computed in $O(n)$ time.*

This result relates to previous studies on topological rectangle visibility graphs [20] and transversal structures (see, e.g., [12,13,17,19]). Also, starting from a Φ -LSVR of $\langle G_r, G_b \rangle$, we can compute a *simultaneous RAC embedding* (SRE) of the two graphs with at most two bends per edge, improving the general upper bound by Bekos *et al.* [3] for those pairs of graphs that can be directed and augmented to admit a Φ -LSVR. As an application of this result, in Section 5 we show an alternative proof of another result by Bekos *et al.* that a wheel graph and a matching admit an SRE with at most two bends for each edge of the wheel, and no bends for the matching edges [3].

The proof of Theorem 1 is based on a characterization described in Section 3, which allows for an efficient testing algorithm presented in Section 4. Section 2 contains preliminaries. In Section 6 we conclude with a discussion on possible research directions that arise from our research.

2 Preliminaries

A graph $G = (V, E)$ is *simple*, if it contains neither loops nor multiple edges. We consider simple graphs, if not otherwise specified. A *drawing* Γ of G maps each vertex of V to a point of the plane and each edge of E to a Jordan arc between its two endpoints. We only consider *simple drawings*, i.e., drawings such that the arcs representing two edges have at most one point in common, which is either a common end-vertex or a common interior point where the two arcs properly cross. A drawing is *planar* if no two arcs representing two edges cross. A planar drawing subdivides the plane into topologically connected regions, called *faces*. The unbounded region is called the *outer face*. A *planar embedding* of a graph is an equivalence class of planar drawings that define the same set of faces. A graph with a given planar embedding is a *plane* graph. For a non-planar drawing, we can still derive an embedding considering that the boundary of a face may consist also of edge segments between vertices and/or crossing points of edges. The unbounded region is still called the outer face.

A graph is *biconnected* if it remains connected after removing any one vertex. A directed graph (a digraph for short) is biconnected if its underlying undirected graph is biconnected. The *dual graph* D of a plane graph G is a plane multigraph whose vertices are the faces of G with an edge between two faces if and only if they share an edge. If G is a digraph, D is also a digraph whose dual edge e^* for a primal edge e is conventionally directed from the face, $left_G(e)$, on the left of e to the face, $right_G(e)$, on the right of e . Since we also use the opposite convention, we let D^\rightarrow (resp., D^\leftarrow) be the dual whose edges cross the primal edges from left to right (resp., right to left).

A *topological numbering* of a digraph is an assignment, X , of numbers to its vertices such that $X(u) < X(v)$ for every edge (u, v) . A graph admits a topological numbering if and only if it is acyclic. An acyclic digraph with a single source s and a single sink t is called an *st-graph*. In such a graph, for every vertex v , there exists a directed path from s to t that contains v [22]. A *plane st-graph* is an *st-graph* that is planar and embedded such that s and t are on the boundary of the outer face. In any *st-graph*, the presence of the edge (s, t) guarantees that the graph is biconnected. In the following we consider *st-graphs* that contain the edge (s, t) , as otherwise it can be added without violating planarity. Let G be a plane *st-graph*, then for each vertex v of G the incoming edges appear consecutively around v , and so do the outgoing edges. Vertex s only has outgoing edges, while vertex t only has incoming edges. This is a particular transversal structure (see Section 3) known as a *bipolar orientation* [18,22]. Each face f of G is bounded by two directed paths with a common origin and destination, called the *left path* and *right path* of f . For all vertices v and edges e on the left (resp., right) path of f , we let $right_G(v) = right_G(e) = f$ (resp., $left_G(v) = left_G(e) = f$).

Tamassia and Tollis [22] proved the following lemma.

Lemma 1 ([22]). *Let G be a plane st -graph and let D^\rightarrow be its dual graph. Let u and v be two vertices of G . Then exactly one of the following four conditions holds:*

1. G has a path from u to v .
2. G has a path from v to u .
3. D^\rightarrow has a path from $\text{right}_G(u)$ to $\text{left}_G(v)$.
4. D^\rightarrow has a path from $\text{right}_G(v)$ to $\text{right}_G(u)$.

Let v be a vertex of G , then denote by $B(v)$ (resp., $T(v)$) the set of vertices that can reach (resp., can be reached from) v . Also, denote by $L(v)$ (resp., $R(v)$) the set of vertices that are to the left (resp., to the right) of every path from s to t through v . By Lemma 1, these four sets partition the vertices of $G \setminus \{v\}$. In every planar drawing of G , they are contained in four distinct regions of the plane that share point v . The vertices of $B(v)$ are in the region delimited by the leftmost and the rightmost paths from s to v , while the vertices of $T(v)$ are in the region delimited by the leftmost and the rightmost paths from v to t . Edge (s, t) separates the two regions containing the vertices of $L(v)$ and $R(v)$, as in Fig. 1. Refer to [8] for further details.




Fig. 1. Vertex sets $B(v)$, $T(v)$, $L(v)$, and $R(v)$ and their corresponding regions of the plane.

3 Characterization

A *transversal structure* of a plane graph G , is a coloring and an orientation of the inner edges (i.e., those edges that do not belong to the outer face) of the graph that obey some local and global conditions. Transversal structures have been widely studied and important applications have been found. Bipolar orientations (also known as *st-orientations*) of plane graphs have been used to compute bar visibility representations [18,22]. Further applications can be found in [12,13,17,19]. See also the survey by Eppstein [11].

To characterize those pairs of graphs that admit a Φ -LSVR, we introduce a new transversal structure for the union of the two graphs (which may be non-planar) and show that it is in bijection with the desired representation. In what follows $G_r = (V_r = V \cup \{s_r, t_r\}, E_r)$ and $G_b = (V_b = V \cup \{s_b, t_b\}, E_b)$ are two plane st -graphs with duals D_r^\rightarrow and D_b^\leftarrow , respectively.

Definition 1. *Given $\Phi : V \rightarrow \mathcal{H} = \{\perp, \lceil, \rceil, \sqcap, \sqcup\}$, a (4-polar) Φ -transversal is a drawing of a directed (multi)graph on the vertex set $V \cup \{s_r, t_r, s_b, t_b\}$ whose edges are partitioned into red edges, blue edges, and the four special edges (s_r, s_b) , (s_b, t_r) , (t_r, t_b) , and (t_b, s_r) forming the outer face, in clockwise order. In addition, the Φ -transversal obeys the following conditions:*

- c1.** *The red (resp., blue) edges induce an st -graph with source s_r (resp., s_b) and sink t_r (resp., t_b).*












Fig. 2. (a) The replacement of the L-shape, ℓ_u , for vertex u with its corner point c_u and the drawing of u 's adjacent edges with 2 bends per edge when constructing a Φ -transversal from a Φ -LSVR. Only ℓ_u 's visibilities are shown. (b) Illustration for the proof of Lemma 3: the case when u is in $B(v)$ and v is in $T(u)$.

- c2.** For every vertex $u \in V$, the clockwise order of the edges incident to u forms four non-empty blocks of monochromatic edges, such that all edges in the same block are either all incoming or all outgoing with respect to u . The four blocks are encountered around u depending on $\Phi(u)$ as in the following table.

 \Rightarrow 	 \Rightarrow 	 \Rightarrow 	 \Rightarrow 
$f_b(u) = \text{right}_b(u)$ $f_r(u) = \text{left}_r(u)$	$f_b(u) = \text{right}_b(u)$ $f_r(u) = \text{right}_r(u)$	$f_b(u) = \text{left}_b(u)$ $f_r(u) = \text{right}_r(u)$	$f_b(u) = \text{left}_b(u)$ $f_r(u) = \text{left}_r(u)$

- c3.** Only blue and red edges may cross and only if blue crosses red from left to right.

A pair of plane st -graphs $\langle G_r, G_b \rangle$ admits a Φ -transversal if there exists a Φ -transversal G_{rb} such that restricting $G_{rb} \setminus \{s_b, t_b\}$ to the red edges realizes the planar embedding G_r and restricting $G_{rb} \setminus \{s_r, t_r\}$ to the blue edges realizes the planar embedding G_b .

Let u be a vertex of V , then the edges of a single color enter and leave u by the same face in the embedding of the other colored graph. In other words, as condition **c2** indicates, $\Phi(u)$ defines the face of G_b (resp., G_r), denoted by $f_b(u)$ (resp., $f_r(u)$), by which the edges of G_r (resp., G_b) incident to u enter and leave u , in the Φ -transversal. Also, condition **c3** implies that edges $\{(s_r, s_b), (s_b, t_r), (t_r, t_b), (t_b, s_r)\}$ are not crossed, because they are not colored.

In the remainder of this section we will prove the next theorem.

Theorem 2. Let G_r and G_b be two plane st -graphs defined on the same set of vertices V and with distinct sources and sinks. Let $\Phi : V \rightarrow \mathcal{H} = \{\text{L}, \text{J}, \text{Gamma}, \text{Gamma}\}$. Then $\langle G_r, G_b \rangle$ admits a Φ -LSVR if and only if it admits a Φ -transversal.

The necessity of the Φ -transversal is easily shown. Let $\langle \Gamma_r, \Gamma_b \rangle$ be a Φ -LSVR of $\langle G_r, G_b \rangle$ with two additional horizontal bars at the bottommost and topmost sides of the drawing that represent s_r and t_r , and two additional vertical bars at the leftmost and rightmost sides of the drawing that represent s_b and t_b . From such a drawing we can extract a Φ -transversal G_{rb} as follows. Since the four vertices s_r, t_r, s_b , and t_b are

represented by the extreme bars in the drawing, these four vertices belong to the outer face, and the four edges on the outer face can be added without crossings. Also, we color red all inner edges represented by vertical visibilities (directed from bottom to top), and blue all inner edges represented by horizontal visibilities (directed from left to right). To see that conditions **c1**, **c2** and **c3** are satisfied, let G_{rb} be a polyline drawing of computed as follows. Let c_u be the corner of the L-shape, ℓ_u , representing vertex u . For every edge (u, v) , replace its visibility segment by a polyline from c_u to c_v that has two bends, both contained in the visibility segment and each at distance δ from a different one of its endpoints, for an arbitrarily small, fixed $\delta > 0$. See Fig. 2(a). Finally, replace every L-shape ℓ_u with its corner c_u . Since each bar visibility representation preserves the embedding of the input graph, **c1** is respected. Also, **c2** and **c3** are clearly satisfied by the embedding derived from G_{rb} . We remark that, by construction, each edge is represented by a polyline with two bends and two edges cross only at right angles; this observation will be used in Section 6.

To prove sufficiency, assume $\langle G_r, G_b \rangle$ admits a Φ -transversal G_{rb} . We present an algorithm, $\Phi\text{LSVRDrawer}$, that takes as input G_{rb} and returns a $\Phi\text{-LSVR}$ $\langle \Gamma_r, \Gamma_b \rangle$ of $\langle G_r, G_b \rangle$. We first recall the algorithm by Tamassia and Tollis (TT in the following) to compute an embedding preserving bar visibility representation of a plane st -graph G , see [8,22]:

1. Compute the dual D^\rightarrow of G .
2. Compute a pair of topological numberings Y of G and X of D^\rightarrow .
3. Draw each vertex v as a horizontal bar with y -coordinate $Y(v)$ and between x -coordinates $X(\text{left}_G(v))$ and $X(\text{right}_G(v)) - \epsilon$.
4. Draw each edge $e = (u, v)$ as a vertical segment at x -coordinate $X(\text{left}_G(e))$, between y -coordinates $Y(u)$ and $Y(v)$, and with thickness ϵ .

We are now ready to describe algorithm $\Phi\text{LSVRDrawer}$.

- Step 1:** Compute the dual graphs D_r^\rightarrow of G_r and D_b^\leftarrow of G_b .
- Step 2:** Compute a pair of topological numberings n_r of G_r and n_b of G_b .
- Step 3:** Compute a pair of topological numberings n_r^* of D_r^\rightarrow and n_b^* of D_b^\leftarrow .
- Step 4:** Compute a bar visibility representation Γ_r of G_r by using the TT algorithm with $X(u) = X_r(u) = n_r^*(u)$ and $Y(u) = Y_r(u) = n_b^*(f_b(u)) + n_r(u)\delta$, for each vertex u . Also, shift the horizontal segment for each vertex u to the left by $n_b(u)\delta$.
- Step 5:** Compute a bar visibility representation Γ'_b of G_b by using the TT algorithm with $X(u) = X_b(u) = n_b^*(u)$ and $Y(u) = Y_b(u) = n_r^*(f_r(u)) + n_b(u)\delta$, for each vertex u . Then turn Γ'_b into a vertical bar visibility representation, Γ_b , by drawing every horizontal segment $((x_0, y), (x_1, y))$ in Γ'_b as the vertical segment $((y, x_0), (y, x_1))$ in Γ_b . Finally, shift the vertical segment for each vertex u up by $n_r(u)\delta$.

Lemma 2 guarantees that Y_r and Y_b are valid topological numberings, and thus, that Γ_r and Γ_b are two bar visibility representations. Also, Lemma 3 ensures the union of Γ_r and Γ_b is a $\Phi\text{-LSVR}$. The shifts performed at the end of **Steps 4-5** are to prevent the bars of two L-shapes from coinciding. The value $\delta > 0$ is chosen to be less than ϵ and less than the smallest difference between distinct numbers divided by the largest number from any topological numbering n_r , n_b , n_r^* , or n_b^* . This choice of δ guarantees that all visibilities are preserved after the shift, and that no new visibilities are introduced.

Lemma 2. Y_r is a valid topological numbering of G_r and Y_b is a valid topological numbering of G_b .

Proof. Let (u, v) be a red edge from u to v . We know that $n_r(u) < n_r(v)$. Let e_0, e_1, \dots, e_k be the blue edges crossed by (u, v) in G_{rb} . Due to conditions **c2** and **c3**, there exists a path $\{f_b(u) = \text{right}_b(e_0), \text{left}_b(e_0) = \text{right}_b(e_1), \dots, \text{left}_b(e_{k-1}) = \text{right}_b(e_k), \text{left}_b(e_k) = f_b(v)\}$ in D_b^- . Thus, we also know that $n_b^*(f_b(u)) \leq n_b^*(f_b(v))$. Since $Y_r(u) = n_b^*(f_b(u)) + n_r(u)\delta$ and $\delta > 0$, it follows that $Y_r(u) < Y_r(v)$. A symmetric argument shows $Y_b(u) < Y_b(v)$ if (u, v) is a blue edge. \square

Lemma 3. Each vertex u of V is represented by an L-shape ℓ_u in $\langle \Gamma_r, \Gamma_b \rangle$ as defined by the function Φ . Also no two L-shapes intersect each other.

Proof. Suppose $\Phi(u) = \text{L}$, as the other cases are similar. Then, $f_b(u) = \text{right}_b(u)$ and $f_r(u) = \text{left}_r(u)$. The horizontal bar representing u in Γ_r is the segment $[p_0(u), p_1(u)]$, where the two points $p_0(u)$ and $p_1(u)$ are $p_0(u) = (n_r^*(\text{left}_r(u)) + n_b(u)\delta, Y_r(u))$, and $p_1(u) = (n_r^*(\text{right}_r(u)) + n_b(u)\delta, Y_r(u))$. Note that, $n_r^*(\text{left}_r(u)) < n_r^*(\text{right}_r(u))$. The vertical bar representing u in Γ_b is the segment $[q_0(u), q_1(u)]$, where the two points $q_0(u)$ and $q_1(u)$ are $q_0(u) = (Y_b(u), n_b^*(\text{right}_b(u)) + n_r(u)\delta)$, and $q_1(u) = (Y_b(u), n_b^*(\text{left}_b(u)) + n_r(u)\delta)$. Note that, $n_b^*(\text{right}_b(u)) < n_b^*(\text{left}_b(u))$. Since $Y_r(u) = n_b^*(f_b(u)) + n_r(u)\delta = n_b^*(\text{right}_b(u)) + n_r(u)\delta$, the bottom coordinate of the vertical bar representing u matches the y -coordinate of the horizontal bar representing u . Since $Y_b(u) = n_r^*(f_r(u)) + n_b(u)\delta = n_r^*(\text{left}_r(u)) + n_b(u)\delta$, the left coordinate of the horizontal bar representing u matches the x -coordinate of the vertical bar representing u . Thus the two bars form the L-shape L .

We now show that no two L-shapes properly intersect each other. Suppose by contradiction that the vertical bar of a vertex u , properly intersects the horizontal bar of a vertex v . Based on Φ , the vertical bar of u involved in the intersection is either a left vertical bar or a right vertical bar, and it is drawn at x -coordinate $n_r^*(\text{left}_r(u)) + n_b(u)\delta$ or $n_r^*(\text{right}_r(u)) + n_b(u)\delta$, respectively. Suppose it is a left vertical bar, as the other case is symmetric. Since u 's vertical bar properly intersects v 's horizontal bar, we know by construction that $n_r^*(\text{left}_r(v)) + n_b(v)\delta < n_r^*(\text{left}_r(u)) + n_b(u)\delta < n_r^*(\text{right}_r(v)) + n_b(v)\delta$. Proper intersection implies that these inequalities are strict, that there is a path in the red dual D_r^- from $\text{left}_r(v)$ to $\text{left}_r(u)$ to $\text{right}_r(v)$, and that the three faces are distinct. This implies that u belongs either to $B_r(v)$ or to $T_r(v)$, and it lies in the corresponding regions of the plane, with $f_r(u)$ (and hence the start/end of curves representing blue edges incident to u) inside the region. Similarly, by considering the blue dual D_b^- , $n_b^*(\text{right}_b(u)) + n_r(u)\delta < n_b^*(f_b(v)) + n_r(v)\delta < n_b^*(\text{left}_b(u)) + n_r(u)\delta$, we know that v belongs either to $B_b(u)$, or to $T_b(u)$, and it lies in the corresponding regions of the plane, with $f_b(v)$ (and hence the start/end of curves representing red edges incident to v) inside the region. No matter which region, $B_r(v)$ or $T_r(v)$, vertex u lies in, or which region, $B_b(u)$ or $T_b(u)$, vertex v lies in, the directed boundary of the blue region ($B_b(u)$ or $T_b(u)$) containing v crosses the directed boundary of the red region ($B_r(v)$ or $T_r(v)$) containing u from right to left. This either violates condition **c3** (if edges of the boundaries cross) or it violates condition **c2** (if the boundaries share a vertex). See Fig. 2(b) for an illustration. \square

Theorem 3. Let G_r and G_b be two plane st -graphs defined on the same set of n vertices V and with distinct sources and sinks. Let $\Phi : V \rightarrow \mathcal{H} = \{\perp, \lrcorner, \top, \Gamma\}$. If $\langle G_r, G_b \rangle$ admits a Φ -transversal, then algorithm $\Phi\text{LSVR}_{\text{Drawer}}$ computes a Φ -LSVR of $\langle G_r, G_b \rangle$ in $O(n)$ time.

Proof. Lemmas 2 and 3 imply that $\Phi\text{LSVR}_{\text{Drawer}}$ computes a Φ -LSVR of $\langle G_r, G_b \rangle$. Computing the dual graphs and the four topological numberings (**Steps 1-3**), as well as computing the two bar visibility representations and shifting each segment (**Steps 4-5**), can be done in $O(n)$ time, as shown in [8,22]. \square

4 Testing Algorithm

In this section, we first show that there exists a pair of plane st -graphs $\langle G_r, G_b \rangle$ that does not admit an Φ -LSVR for any possible function Φ . This emphasizes the need for an efficient testing algorithm. Then, we show how to test whether two plane st -graphs with the same set of vertices admit a Φ -LSVR for a given function ϕ .

Theorem 4. There exists a pair of plane st -graphs $\langle G_r, G_b \rangle$ that does not admit a Φ -LSVR for any possible function Φ .

Proof. Let G_r be the plane st -graph drawn red in Fig. 3 (observe that it is a series-parallel graph, i.e., a partial 2-tree). Also, let G_b be any plane st -graph containing the blue edge (u, v) in Fig. 3.




Fig. 3. Illustration for the proof of Theorem 4.

Edge (u, v) is incident to either $right_r(u)$ or $left_r(u)$ and to either $right_r(v)$ or $left_r(v)$, based on Φ . In any case edge (u, v) will cross a red edge from left to right and a red edge from right to left, which implies that $\langle G_r, G_b \rangle$ does not admit any Φ -transversal, since condition **c3** cannot be respected. It follows that, by Theorem 2, $\langle G_r, G_b \rangle$ does not admit a Φ -LSVR. \square

Our testing algorithm exploits the interplay between the primal of the blue (red) graph and the dual of the red (blue) graph. Given the circular order of the edges around each vertex imposed by the function ϕ , we aim to compute a suitable path in the red dual for each blue edge. Such paths will then be used to route the blue edges. Finally, we check that no two blue edges cross.

We first introduce a few definitions. Let G and D^\rightarrow be a plane st -graph and its dual. Let f and g be two faces of G that share an edge $e = (x, z)$ of G , such that e belongs to the right (resp., left) path of f (resp., g). Let e^* be the dual edge in D^\rightarrow corresponding to e . Let w be a vertex on the right path of f (or, equivalently, on the left path of g). Then w is *cut from above* (resp., *below*) by e^* , if w precedes z (resp., succeeds x) along the right path of f , i.e., all vertices that precede z (including x) are cut from above, while all vertices that succeed x (including z) are cut from below by e^* .

Let G_r and G_b be a pair of plane st -graphs with the same vertex set V and with distinct sources and sinks. Let $\Phi : V \rightarrow \mathcal{H} = \{\perp, \lrcorner, \top, \Gamma\}$. Recall that, for a given vertex u of G_b , with the notation $L_b(u)$, $R_b(u)$, $T_b(u)$ and $B_b(u)$ we represent the set of vertices to the left, to the right, that are reachable from, and that can reach u in G_b , respectively (see Section 2). Then consider an edge $e = (u, v)$ of G_b and a path³ $\pi_e = \{f_r(u) = f_0, e_0^*, f_1, \dots, f_{k-1}, e_{k-1}^*, f_r(v) = f_k\}$ in D_r^\rightarrow , where f_i ($0 \leq i \leq k$) are the faces traversed by the path, and e_i^* ($0 \leq i < k$) are the dual edges used by the path to go from f_i to f_{i+1} . Path π_e is a *traversing path* for e , if $\pi_e = \{f_r(u) = f_r(v)\}$, or for all $0 \leq i < k$ and all vertices w in the right path of f_i :

- p1.** If $w \in L_b(u)$ then w is cut from below by e_i^* . See Fig. 4(a).
- p2.** If $w \in R_b(u)$ then w is cut from above by e_i^* . See Fig. 4(b).
- p3.** If $w \in B_b(u)$ and $\Phi(w) = \perp$ (resp., $\Phi(w) = \Gamma$) then w is cut from above (resp., below) by e_i^* . See Fig. 4(c).
- p4.** If $w \in T_b(u)$ and $\Phi(w) = \Gamma$ (resp., $\Phi(w) = \perp$) then w is cut from above (resp., below) by e_i^* . See Fig. 4(d).
- p5.** If $w \in B_b(v)$ and $\Phi(w) = \top$ (resp., $\Phi(w) = \lrcorner$) then w is cut from above (resp., below) by e_i^* . See Fig. 4(e).
- p6.** If $w \in T_b(v)$ and $\Phi(w) = \lrcorner$ (resp., $\Phi(w) = \top$) then w is cut from above (resp., below) by e_i^* . See Fig. 4(f).

We now show that if $\langle G_r, G_b \rangle$ admits a Φ -transversal, then for each blue edge (the same argument would apply for red edges) there exists a unique traversing path.

Lemma 4. *Let G_r and G_b be two plane st -graphs with the same vertex set V and with distinct sources and sinks. Let $\Phi : V \rightarrow \mathcal{H} = \{\perp, \lrcorner, \top, \Gamma\}$. If $\langle G_r, G_b \rangle$ admits a Φ -transversal, then for every edge e of G_b there is a unique traversing path π_e in D_r^\rightarrow .*

Proof. If $\langle G_r, G_b \rangle$ admits a Φ -transversal G_{rb} , then for every edge $e = (u, v)$ of G_b there exists a path $\pi_e = \{f_r(u) = f_0, e_0^*, f_1, \dots, f_{k-1}, e_{k-1}^*, f_r(v) = f_k\}$ in D_r^\rightarrow , which is the path used by e to go from $f_r(u)$ to $f_r(v)$ in G_{rb} .

If f_0 and f_k coincide, then π_e is a traversing path. Otherwise, we would have a cycle $\pi_e = \{f_0 = f_k, \dots, f_0 = f_k\}$, which is not possible since D_r^\rightarrow is acyclic, being the dual of a plane st -graph.

If f_0 and f_k do not coincide, let w be a vertex in the right path of f_i . First, if w belongs to $L_b(u)$, then it is cut from below. Otherwise, if w was cut from above, since edge $e = (u, v)$ cannot cross the right path of f_i twice (by condition **c3**), it would belong to $R_b(u)$, a contradiction with the fact that the embedding of G_b is preserved. Thus **p1**

³ Since D_r^\rightarrow is a multigraph, to uniquely identify π_e we specify the edges that are traversed.




Fig. 4. Illustration for the properties of a traversing path π_e .

is respected by π_e . With a symmetric argument we can also prove **p2**. Suppose now that w belongs to $B_b(u)$, then $f_r(w) = f_i = \text{left}_r(w)$, otherwise if $f_r(w) = f_{i+1} = \text{right}_r(w)$, the blue path from w to u would violate **c3**. In other words, either $\Phi(w) = \sqsubset$ or $\Phi(w) = \sqsupset$. Furthermore, if $\Phi(w) = \sqsubset$, then w must be cut from above, while if $\Phi(w) = \sqsupset$, then w must be cut from below, as otherwise the incoming blue edges to w must enter a region delimited by the blue path from w to u , the blue edge (u, v) , and part of the (red) right path of f_i , which violates the planarity of the embedding of G_b or condition **c2** (see Fig. 4(c)). Thus **p3** is respected by π_e . With similar arguments one can prove **p4** – **p6**. Hence, π_e is a traversing set. To prove that π_e is unique, note that any possible traversing set for e must start from f_0 and leave this face. Hence, any vertex w on the right path of f_0 must be cut from either above or below, according to properties **p1** – **p6** (which cover all possible cases for w). The only edge that can satisfy the cut condition for all vertices on the right path of f_0 , is an edge e_0^* whose corresponding red primal edge, denoted by (x, z) , is such that all vertices on the right path of f_0 above x must be cut from below and all those below z must be cut from above. Clearly, this edge is unique. By repeatedly applying this argument for each face f_i ($0 \leq i < k$), the traversing path π_e is uniquely identified. \square

The next theorem concludes the proof of Theorem 1.

Theorem 5. Let G_r and G_b be two plane st-graphs with the same set of n vertices V and with distinct sources and sinks. Let $\Phi : V \rightarrow \mathcal{H} = \{\sqsubset, \sqsupset, \sqcap, \sqcup\}$. There exists an $O(n^3)$ -time algorithm to test whether $\langle G_r, G_b \rangle$ admits a Φ -transversal.

Proof. Our testing algorithm aims to compute (if it exists) a Φ -transversal G_{rb} for $\langle G_r, G_b \rangle$. We first fix the circular order of the edges restricted to the blue edges (resp., red edges) around each vertex u of G_{rb} to satisfy **c1** and to maintain the planar embedding of G_b (resp., G_r). We then fix the circular order of the blue edges with respect to the red edges around each vertex u of G_{rb} to satisfy **c2** (i.e., to obey $\Phi(u)$). Then, we first check if for every blue edge e there exists a traversing path π_e ; if so, we verify that by routing every blue edge e through π_e no two blue edges cross each other. If this procedure succeeds then $\langle G_r, G_b \rangle$ admits Φ -transversal G_{rb} , because, by construction, the resulting embedding of G_{rb} satisfies conditions **c1**, **c2** and **c3**. Otherwise, either there exists a blue edge with no traversing path, or two traversing paths are such that the two corresponding edges of G_b cross if routed through them. In the first case $\langle G_r, G_b \rangle$ does not admit a Φ -transversal by Lemma 4. In the second case, since the traversing paths are unique, condition **c2** cannot be satisfied, and again $\langle G_r, G_b \rangle$ does not admit a Φ -transversal.

The testing algorithm works in two phases as follows.

Phase 1. For every edge $e = (u, v) \in E_b$. If $f_r(u) = f_r(v)$, we have found a traversing path. Otherwise, we label each vertex on the right path of $f_r(u)$, by A if it must be cut from above or by B if it must be cut from below, according to properties **p1** – **p6**. Then we check if the sequence of labels along the path is a nonzero number of A 's followed by a nonzero number of B 's. If so, then the dual edge of the traversing path is the one whose corresponding primal edge has the two end-vertices with different labels (which is unique). If this is not the case, then a traversing path for e does not exist. In the positive case, we add the dual edge we found and the next face we reach through this edge to π_e and we iterate the algorithm until we reach either $f_r(v)$ or the outer face of D_r^\rightarrow . In the former case π_e is a traversing path for e , while in the latter case, since the edges of the outer face of G_{rb} cannot be crossed by definition of Φ -transversal, we have that again no traversing path can be found.

Phase 2. We now check that by routing every edge $e \in E_b$ through its corresponding traversing path π_e , no two of these edges cross each other. Consider the dual graph D_r^\rightarrow , which is a plane st -graph. Construct a planar drawing Γ of D_r^\rightarrow . Consider any two traversing paths π_e and $\pi_{e'}$, which corresponds to two paths in Γ , and let $e = (u, v)$ and $e' = (w, z)$ be the two corresponding edges of G_b . Denote by $\hat{\pi}_e = \{u\} \cup \pi_e \cup \{v\}$ and $\hat{\pi}_{e'} = \{w\} \cup \pi_{e'} \cup \{z\}$ the two enriched paths. Enrich Γ by adding the four edges $(x, f_r(x))$, where $x \in \{u, v, w, z\}$, in a planar way respecting the original embedding of G_b . Consider now the subdrawing Γ' of Γ induced by $\hat{\pi}_e \cup \hat{\pi}_{e'}$. See Fig. 5 for an illustration. If e and e' cross each other, then $\pi_e \cap \pi_{e'}$ cannot be empty. Moreover, the intersection $\pi_e \cap \pi_{e'}$ must be a single subpath, as otherwise the two traversing paths would not be unique. Let f be the first face and let g be the last face in this subpath. Let e_u be the incoming edge of f that belongs to the subpath of $\hat{\pi}_e$ from u to f ; and let e_w be the incoming edge of f that belongs to the subpath of $\hat{\pi}_{e'}$ from w to f . Also, let e_v be the outgoing edge of g that belongs to the subpath of $\hat{\pi}_e$ from g to v ; and let e_z be the outgoing edge of g that belongs to the subpath of $\hat{\pi}_{e'}$ from g to z . Then e and e' cross if and only if walking clockwise along $\pi_e \cup \pi_{e'}$ from f to g and back to f these four edges are encountered in the circular order e_u, e_z, e_v, e_w as shown in Fig. 5. Note that, e_u and e_w may coincide if $u = w$, and similarly for e_v and e_z .




Fig. 5. Illustration for the proof of Theorem 5.

To conclude the proof we need to show that **Phase 1** and **Phase 2** can be implemented in $O(n^3)$ time. Checking if a traversing path exists for each edge in E_b (**Phase 1**) can be performed in $O(n^3)$ time. Namely, for each of these edges, we need to assign a label A or B to every vertex w on the right path of the traversed faces. This can be done by checking the existence of a path in the primal graph G_b or in the dual graph D_b^+ , as by Lemma 1, thus requiring $O(n)$ time per vertex, $O(n^2)$ time per edge, and $O(n^3)$ time in total. Furthermore, checking if two of these edges would cross if routed through their traversing paths (**Phase 2**) can also be performed in $O(n^3)$ time. In fact, constructing a planar drawing of D_r^+ costs $O(n)$ time [7], while checking if two paths cross in the drawing costs $O(n)$ time for each pair of paths, of which there are $O(n^2)$ in total. This concludes the proof. \square

5 Simultaneous RAC Embeddings via Simultaneous Visibility Representations with L-shapes

In this section, we show that Theorem 1 can be used to shed more light on the problem of computing a simultaneous RAC embedding [2,3].

Given two planar graphs with the same vertex set, an SRE is a simultaneous embedding where crossings between edges of the two graphs occur at right angles. Argyriou *et al.* proved that it is always possible to construct an SRE with straight-line edges of a cycle and a matching, while there exist a wheel graph and a cycle that do not admit such a representation [2]. This motivated recent results about SRE with bends along the edges. Namely, Bekos *et al.* show that two planar graphs with the same vertex set admit an SRE with at most six bends per edge in both graphs [3].

We observe that any pair of graphs that admit a simultaneous visibility representation with L-shapes also admits an SRE with at most two bends per edge. This is obtained with the technique used in Section 3 to compute a Φ -transversal from a Φ -LSVR, see Fig. 2(a). Thus, a new approach to characterize graph pairs that have SREs with at most two bends per edge is as follows: Given two planar graphs with the same vertex set, add




Fig. 6. The wheel W_{10} , edges of a matching M , and the Φ -transversal representing both.

to each of them a unique source and a unique sink, and look for two st -orientations (one for each of the two graphs) and a function Φ such that the two graphs admit a Φ -LSVR.

As an application of this result, in what follows we show an alternative proof of another result by Bekos *et al.* that a wheel graph and a matching admit an SRE with at most two bends for each edge of the wheel, and no bends for the matching edges [3].

The *wheel* on n vertices, denoted by W_n , is an undirected graph with vertices $V = \{1, 2, \dots, n\}$ and edges $\{(i, n) : 1 \leq i \leq n-1\}$, called *spokes*; and $\{(i, i+1) : 1 \leq i \leq n-2\} \cup \{(1, n-1)\}$, called *rim edges*. See Fig. 6 for an illustration. A graph is a *matching* if no two of its edges share a common end-vertex and a *complete matching* if every vertex is the end-vertex of some edge.

We begin with the following observation, based on Theorem 2.

Observation 1. Let $\langle G'_r = (V, E'_r), G'_b = (V, E'_b) \rangle$ be a pair of undirected plane graphs. Let $\Phi : V \rightarrow \mathcal{H} = \{\perp\}$. Then $\langle G'_r, G'_b \rangle$ admits a Φ -LSVR, if and only if G'_r and G'_b can be augmented and st -oriented to two plane st -graphs G_r and G_b , respectively, defined on the same set of vertices and with distinct sources and sinks, such that $\langle G_r, G_b \rangle$ admits a Φ -transversal.

We show that the wheel W_n and any matching M on the vertices of W_n has a Φ -LSVR where $\Phi(v) = \perp$ for all vertices v in W_n . To do this, we construct a Φ -transversal for the pair $\langle G_r, G_b \rangle$, where G_r and G_b are plane st -graphs obtained from W_n and M , respectively.

We assume that n is even and M is a complete matching, though we will later see how to handle partial matchings. By symmetry, we may assume that vertex n is matched with vertex $n-1$. Let $e_1, e_2, \dots, e_{n/2}$ be the edges in M where $e_1 = (n-1, n)$. Let

the *height*, $h(u)$, of a vertex u be the index of its edge in the matching, so if $u \in e_i$ then $h(u) = i$.

For all $u \in \{1, 2, \dots, n-2\}$, use the point $(u, h(u))$ to represent vertex u . Use the point $(n/2, -n^2/2)$ to represent vertex n . Use the point $(n/2, -n^2/2 - 1)$ to represent vertex $n-1$. The drawing of W_n is straight-line planar because the spokes of the wheel do not intersect the rim edges since the slope of any rim edge (excluding $(n-2, n-1)$ and $(1, n-1)$) is in the range $(-n/2, +n/2)$.

To define G_r (the *st*-graph representing W_n), direct every rim edge (u, v) from smaller to larger height vertex. If both vertices have the same height (i.e., the edge is part of the matching), direct the edge from smaller to larger numbered vertex. Direct spoke edges (n, u) for $1 \leq u \leq n-2$ from vertex n to vertex u . Direct spoke edge $(n-1, n)$ from vertex $n-1$ to vertex n . Add the directed edge $(s_r, n-1)$ and the directed edges (u, t_r) for all $1 \leq u \leq n-1$, where the points $(n/2, -\infty)$ and $(n/2, +\infty)$ represent s_r and t_r respectively. Finally, add the edge (s_r, t_r) on the outer face of G_r .

The drawing of M is also straight-line planar since its edges are all horizontal segments with different positive y -coordinates, except for $(n-1, n)$, which is vertical and spans only negative y -coordinates. To define G_b (the *st*-graph representing M), direct all edges from smaller to larger numbered end-point, so the edges are directed from left to right. Add the directed edge (s_b, u) (resp., (u, t_b)) for each vertex u that is the smaller (resp., larger) numbered end-point of an edge in M , where the points $(-\infty, 0)$ and $(+\infty, 0)$ represent s_b and t_b , respectively. Finally, add the edge (s_b, t_b) on the outer face of G_b .

To ensure condition **c2** and preserve condition **c3** of Theorem 2, we slightly modify the drawing of the edges in G_b (i.e., the blue edges). When drawn as straight-line segments, the head of every blue edge terminates in the left red face of its destination as desired, but the tail does not leave from the left red face of its origin, in general. This is easily fixed by drawing each blue edge (except those incident to s_b) so that it leaves from the origin's left red face, heading to the left and slightly above the entering blue edge, and then loops to the right over the vertex to rejoin its original rightward horizontal path. See Fig. 6.

All red edges (from W_n) are directed from bottom to top. Blue edges (from M), except for $(n-1, n)$ which is identical to the corresponding red edge (from W_n), intersect red edges only during the blue edges' rightward trajectories. Thus blue edges only cross red edges from left to right.

To represent a partial matching P , which is a complete matching M without some edges, for all edges $(u, v) \in M \setminus P$, we add edge (s_b, v) and (u, t_b) to G_b , routing the first slightly above and the second slightly below the existing blue path $\langle s_b, u, v, t_b \rangle$, and remove the edge (u, v) from G_b .

6 Final Remarks and Open Problems

In this paper we have introduced and studied the concept of simultaneous visibility representation with L-shapes of two plane *st*-graphs. We remark that it is possible to include in our theory the case when the vertices can also be drawn as rectangles. Nevertheless, this would not enlarge the class of representable pairs of graphs. In fact, for

every vertex v drawn as a rectangle \mathcal{R}_v , we can replace \mathcal{R}_v with any L-shape by keeping only two adjacent sides of \mathcal{R}_v in the drawing and prolonging the visibilities incident to the removed sides of \mathcal{R}_v . The converse is not true. Indeed, roughly speaking, L-shapes can be nested, whereas rectangles cannot. To give an example, if a vertex v must see a vertex u both vertically and horizontally, this immediately implies that the two corresponding rectangles need to overlap, while two L-shapes could instead be nested. Several extensions of the model introduced in this paper can also be studied, e.g., the case where every edge is represented by a T-shape, or more generally by a $+$ -shape.

Three questions that stem from this paper are whether the time complexity of the testing algorithm in Section 4 can be improved; what is the complexity of deciding if two given plane st -graphs admit a Φ -LSVR for some function Φ , which is not part of the input; and what is the complexity of deciding if two undirected graphs admit a Φ -LSVR for some function Φ .

References

1. P. Angelini, M. Geyer, M. Kaufmann, and D. Neuwirth. On a tree and a path with no geometric simultaneous embedding. *JGAA*, 16(1):37–83, 2012.
2. E. N. Argyriou, M. A. Bekos, M. Kaufmann, and A. Symvonis. Geometric RAC simultaneous drawings of graphs. *JGAA*, 17(1):11–34, 2013.
3. M. A. Bekos, T. C. van Dijk, P. Kindermann, and A. Wolff. Simultaneous drawing of planar graphs with right-angle crossings and few bends. In M. S. Rahman and E. Tomita, editors, *WALCOM 2015*, volume 8973 of *LNCS*, pages 222–233. Springer, 2015.
4. T. Bläsius and I. Rutter. Simultaneous pq-ordering with applications to constrained embedding problems. In S. Khanna, editor, *SODA 2013*, pages 1030–1043. SIAM, 2013.
5. P. Braß, E. Cenek, C. A. Duncan, A. Efrat, C. Erten, D. Ismailescu, S. G. Kobourov, A. Lubiw, and J. S. B. Mitchell. On simultaneous planar graph embeddings. *Comput. Geom.*, 36(2):117–130, 2007.
6. S. Cabello, M. J. van Kreveld, G. Liotta, H. Meijer, B. Speckmann, and K. Verbeek. Geometric simultaneous embeddings of a graph and a matching. *JGAA*, 15(1):79–96, 2011.
7. M. Chrobak and T. H. Payne. A linear-time algorithm for drawing a planar graph on a grid. *Inform. Process. Lett.*, 54(4):241–246, 1995.
8. G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing*. Prentice Hall, 1999.
9. E. Di Giacomo and G. Liotta. Simultaneous embedding of outerplanar graphs, paths, and cycles. *Int. J. Comput. Geometry Appl.*, 17(2):139–160, 2007.
10. W. Didimo and G. Liotta. The crossing angle resolution in graph drawing. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*. Springer, 2012.
11. D. Eppstein. Regular labelings and geometric structures. In *CCCG 2010*, pages 125–130, 2010.
12. S. Felsner. Rectangle and square representations of planar graphs. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 213–248. Springer, 2013.
13. É. Fusy. Transversal structures on triangulations: A combinatorial study and straight-line drawings. *Discr. Math.*, 309(7):1870–1894, 2009.
14. W. Huang, S.-H. Hong, and P. Eades. Effects of crossing angles. In *PacificVis 2008*, pages 41–46. IEEE, 2008.
15. K. R. Jampani and A. Lubiw. Simultaneous interval graphs. In O. Cheong, K. Chwa, and K. Park, editors, *ISAAC 2010*, volume 6506 of *LNCS*, pages 206–217. Springer, 2010.

16. K. R. Jampani and A. Lubiw. The simultaneous representation problem for chordal, comparability and permutation graphs. *JGAA*, 16(2):283–315, 2012.
17. G. Kant and X. He. Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *Theor. Comput. Sci.*, 172(1-2):175–193, 1997.
18. P. Rosenstiehl and R. E. Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discr. & Comput. Geom.*, 1:343–353, 1986.
19. W. Schnyder. Embedding planar graphs on the grid. In D. S. Johnson, editor, *SODA 1990*, pages 138–148. SIAM, 1990.
20. I. Streinu and S. Whitesides. Rectangle visibility graphs: Characterization, construction, and compaction. In H. Alt and M. Habib, editors, *STACS 2003*, volume 2607 of *LNCS*, pages 26–37. Springer, 2003.
21. R. Tamassia. *Handbook of Graph Drawing and Visualization*, chapter Simultaneous embedding of planar graphs. CRC press, 2013.
22. R. Tamassia and I. G. Tollis. A unified approach to visibility representations of planar graphs. *Discr. & Comput. Geom.*, 1(1):321–341, 1986.